

NORMALIZING TOPOLOGICALLY MINIMAL SURFACES I: GLOBAL TO LOCAL INDEX

DAVID BACHMAN

ABSTRACT. We show that in any triangulated 3-manifold, every index n topologically minimal surface can be transformed to a surface which has local indices (as computed in each tetrahedron) that sum to at most n . This generalizes classical theorems of Kneser and Haken, and more recent theorems of Rubinstein and Stocking, and is the first step in a program to show that every topologically minimal surface has a normal form with respect to any triangulation.

1. INTRODUCTION

Minimal surfaces are classical objects of geometry, defined by the condition that their mean curvature is zero everywhere. Equivalently, minimal surfaces represent critical values for surface area, among the space of all surfaces embedded in a given 3-manifold. The word *minimal* is a bit of a misnomer, since such critical values are not necessarily local minima.

Recently, the author introduced *topologically minimal* surfaces [Bac10]. Although their definition is quite different, these surfaces possess many of the same properties as the geometrically minimal surfaces described above. For example, it is well known that a geometrically minimal surface and an incompressible surface can be isotoped so that their intersection curves are essential on both. This has also been shown to be true of topologically minimal surfaces [Bac10].

In recent work, Colding and Minicozzi [CM04a], [CM04b], [CM04c], [CM04d] have classified the possible local pictures of a geometrically minimal surface. For example, given a ball B and a geometrically minimal surface H , Colding and Minicozzi showed that if $H \cap B$ is simply connected, then it is either the graph of a function, or a helicoid. In [Bac10], the author conjectured that every topologically minimal surface can be isotoped to be geometrically minimal. An immediate consequence would be that every topologically minimal surface can be

isotoped so that its possible local pictures are described by the results of Colding and Minicozzi.

This is the first in a series of papers where we study the way in which a topologically minimal surface H can be isotoped with respect to the tetrahedra of a triangulation. Our goal is to show that H can be isotoped so that it meets each tetrahedron in precisely the same way that a geometrically minimal surface can meet a ball, as described by Colding and Minicozzi. In such a position, for example, if Δ is a tetrahedron and $H \cap \Delta$ is simply connected, then we show that $H \cap \Delta$ is contained in either a plane or a helicoid.

1.1. Generalizing normal and almost normal surfaces.

The present work can be viewed from two different perspectives. As described above, it can be thought of as further strengthening the tie between geometrically and topologically minimal surfaces. Alternately, this work can be seen as a natural extension of results of Kneser, Haken, Rubinstein and Stocking on normal and almost normal surfaces. Just as in the case of geometrically minimal surfaces, topologically minimal surfaces have a well defined *index*. Index 0 surfaces are precisely the incompressible surfaces of Haken [Hak61], and index 1 surfaces are the strongly irreducible surfaces of Casson and Gordon [CG87].

In 1929, Kneser [Kne29] showed that an essential sphere S can be isotoped to a normal form with respect to an arbitrary triangulation. In this position, S meets each tetrahedron Δ in the same way a plane might. In particular, $S \cap \Delta$ is a collection of triangles and quadrilaterals, called *normal disks*. In later work, Haken showed that incompressible (i.e. index 0) surfaces can be isotoped into the same normal form [Hak68]. In [Rub95] and [Sto00], Rubinstein and Stocking independently showed that strongly irreducible (i.e. index 1) surfaces can be isotoped to meet each tetrahedron of a given triangulation in a collection of normal disks, with at most one possible exception in one tetrahedron. This exceptional piece is either a saddle-shaped octagon, or two normal disks connected by an unknotted tube (similar to a *catenoid*, a classical geometrically minimal surface).

In [Bac02], the author introduced *critical surfaces*, which are now recognized as the topologically minimal surfaces with index 2 [Bac10]. In unpublished work, both the author [Bac] and Johnson [Joh] claimed that critical surfaces can be isotoped to meet the tetrahedra of a triangulation in a collection of planar pieces, with at most two exceptions. If there are exactly two exceptional pieces, then each is either a tube or an octagon, as described above. If there is a single exceptional piece then it may be a tube or an octagon as well, or something a bit more

complicated. The possibilities are a helical shaped 12-gon, three normal disks connected by two unknotted tubes, an octagon tubed to a normal disk, or an octagon tubed to itself.

In this paper we give a local definition of the index of a surface. This is an index that is computed separately in each tetrahedron. It follows from this definition that the surfaces in a tetrahedron that have index 0 are the normal disks, and the surfaces with index 1 are tubes and octagons. In [Bacc] we show that the surfaces in a tetrahedron with index 2 are the other types of exceptional pieces mentioned above. Using this language, we can concisely summarize all of the results mentioned above as follows:

Theorem 1.1. [Kne29], [Hak68], [Rub95], [Sto00], [Baca], [Joh] *A closed, topologically minimal surface H whose index is $n \leq 2$ in a triangulated 3-manifold can be isotoped so that the sum, taken over all tetrahedra Δ , of the indices of $H \cap \Delta$ is at most n .*

The main result of this paper is a generalization of Theorem 1.1 to arbitrary values of the index n , and to surfaces with possibly non-empty boundary. In particular, we show:

Theorem 1.2. *A properly embedded, topologically minimal surface H whose index is n in a triangulated 3-manifold can be compressed, ∂ -compressed and isotoped to a surface H' so that the sum, taken over all tetrahedra Δ , of the indices of $H' \cap \Delta$ is at most n .*

The compressions in the above theorem are very restricted, and only happen when the initial index n is at least three. In particular, when H is closed and $n \leq 2$, then we recover Theorem 1.1.

Theorem 1.2 is the first step in producing a normal form for any index n topologically minimal surface, with respect to an arbitrary triangulation. To complete the picture of such a normal form, one would have to classify all possible surfaces *in a single tetrahedron* whose index is n . In the first sequel to this paper [Bacb], we begin such a classification. We show there that if a topologically minimal surface (with index at least 1) in a tetrahedron is simply connected then it is a helicoid. Furthermore, the number of “turns” of such a helicoid is directly proportional to its index. This, combined with Theorem 1.2, gives us the desired topological version of the Colding-Minicozzi result mentioned above: any topologically minimal surface in a triangulated 3-manifold can be transformed so that if the intersection with a given tetrahedron is simply connected, then it is either contained in a plane or a helicoid.

In the second sequel to this paper [Bacc] we classify all index 2, non-simply connected surfaces in a tetrahedron. This result, combined with Theorem 1.2 and the results from [Bacb], gives a normal form for all index 2 topologically minimal surfaces in a 3-manifold, with respect to an arbitrary triangulation. Having such a normal form has several known and potential applications. For example, in [Bacc] we combine this normalization result with Corollary 11.2 of [Bacd] to show that there exists a single 3-manifold with infinitely many Heegaard splittings with Hempel distance (see [Hem01]) exactly one, a fact that was not previously known. In joint work with Ryan Derby-Talbot and Eric Sedgwick, we use the normalization of index 2 surfaces to show that the set of Heegaard surfaces of 3-manifolds with toroidal boundary does not change after “generic” Dehn filling [BDTSb]. A potential application is an algorithm to determine if two arbitrary Heegaard surfaces are isotopic, and if they are not, the number of stabilizations required to make them isotopic. A second potential application is an algorithm to determine if a given Heegaard surface is strongly irreducible.

The author thanks Ryan Derby-Talbot and Eric Sedgwick for countless helpful conversations about the present work. Much of the difficulty in working in this area is in finding good notation for the myriad of types of disks and complexes that arise. Most of the notation used here was developed jointly with Derby-Talbot and Sedgwick for the paper [BDTSa], where we present a new proof of the normalization of index 1 surfaces. Many similar ideas are used in both that proof and the present paper.

2. COMPRESSING DISKS AND THE DISK COMPLEX

Throughout this paper H will represent a connected, orientable surface that is properly embedded in a compact, orientable, irreducible 3-manifold M with incompressible boundary and triangulation \mathcal{T} . Furthermore, we will assume H separates M and is not contained in a ball. In this section we give basic definitions and facts regarding compressing disks and associated disk complexes.

Definition 2.1. An embedded loop or arc α on H is said to be *inessential* if it cuts off a subdisk of H , and *essential* otherwise.

We define three types of compressing disks, extending the standard definitions to account for the presence of a given complex \mathcal{K} . We are concerned here with the cases when $\mathcal{K} = \emptyset$, \mathcal{T}^1 , or \mathcal{T}^2 , the 1- and 2-skeleton of our triangulation \mathcal{T} . However, the case when \mathcal{K} is a knot or link in a 3-manifold may be of interest to those who study “thin

position” and related 3-manifold techniques. In what follows, $N(\mathcal{K})$ refers to a regular neighborhood of \mathcal{K} in M .

Definition 2.2. A *compressing disk* C for (H, \mathcal{K}) is a disk embedded in $M - N(\mathcal{K})$ so that $\alpha = C \cap H = \partial C$ is an essential loop in $H - N(\mathcal{K})$.

Let $\mathcal{C}(H, \mathcal{K})$ denote the set of compressing disks for (H, \mathcal{K}) .

Definition 2.3. A ∂ -*compressing disk* B for (H, \mathcal{K}) is a disk embedded in $M - N(\mathcal{K})$ so that $\partial B = \alpha \cup \beta$, where $\alpha = B \cap H$ is an essential arc on $H - N(\mathcal{K})$ and $B \cap \partial M = \beta$. A ∂ -compressing disk is *fake* if the arc β is parallel into ∂H on $\partial M - N(\mathcal{K})$ and *real* otherwise.

Let $\mathcal{B}(H, \mathcal{K})$ denote the set of ∂ -compressing disks for (H, \mathcal{K}) , and $\mathcal{B}^r(H, \mathcal{K})$ the real ∂ -compressing disks.

Definition 2.4. When $\mathcal{K} = \mathcal{T}^1$ or \mathcal{T}^2 , we define an *edge-compressing disk* E for (H, \mathcal{K}) to be a disk embedded in M so that $\partial E = \alpha \cup \beta$, where $\alpha = E \cap H$ is an arc on H and $\beta = E \cap \mathcal{K} \subset e$, for some edge e of \mathcal{T}^1 . We say E is a *boundary* edge-compressing disk when the edge of \mathcal{T}^1 that it is incident to is contained in ∂M , and is an *interior* edge-compressing disk otherwise.

Let $\mathcal{E}(H, \mathcal{K})$ denote the set of edge compressing disks for (H, \mathcal{K}) .

Notice that each element of $\mathcal{E}(H, \mathcal{T}^1)$ meets $M - N(\mathcal{T}^1)$ in a real ∂ -compressing disk for $(H - N(\mathcal{T}^1), \emptyset)$. Hence, $\mathcal{E}(H, \mathcal{T}^1)$ can be identified with a subset of $\mathcal{B}^r(H - N(\mathcal{T}^1), \emptyset)$. This fact will become important in Section 5.

Definition 2.5. Henceforth, we will employ the following shortened notation:

- $\mathcal{C}(H) = \mathcal{C}(H, \emptyset)$
- $\mathcal{B}^r(H) = \mathcal{B}^r(H, \emptyset)$
- $\mathcal{CB}^r(H, \mathcal{K}) = \mathcal{C}(H, \mathcal{K}) \cup \mathcal{B}^r(H, \mathcal{K})$
- $\mathcal{CE}(H, \mathcal{K}) = \mathcal{C}(H, \mathcal{K}) \cup \mathcal{E}(H, \mathcal{K})$
- $\mathcal{D}(H, \mathcal{K}) =$ any one of $\mathcal{C}(H, \mathcal{K})$, $\mathcal{B}^r(H, \mathcal{K})$, $\mathcal{E}(H, \mathcal{K})$, $\mathcal{CB}^r(H, \mathcal{K})$ or $\mathcal{CE}(H, \mathcal{K})$.

Definition 2.6. Suppose $D \in \mathcal{D}(H, \mathcal{K})$. We construct a surface H/D , which is said to have been obtained by *surgering* along D , as follows. Let $M(H, \mathcal{K})$ denote the manifold obtained from $M - N(\mathcal{K})$ by cutting open along H . Let $N(D)$ denote a neighborhood of D in $M(H, \mathcal{K})$. Construct the surface H' in $M - N(\mathcal{K})$ by removing $N(D) \cap H$ from H and replacing it with the frontier of $N(D)$ in $M(H, \mathcal{K})$. The surface H/D is then obtained from H' by discarding any component that lies in a ball in M . See, for example, Figure 1.




FIGURE 1. Surgering along an edge-compressing disk $E \in \mathcal{E}(H, \mathcal{T}^1)$ produces a surface in $M - N(\mathcal{T}^1)$.

Definition 2.7. We say H is *compressible* if $\mathcal{C}(H)$ is non-empty, and *incompressible* otherwise. If D is a compressing disk for H then H/D is said to have been obtained from H by *compressing along* D .

Definition 2.8. Suppose \mathcal{D} is a collection of elements of $\mathcal{D}(H, \mathcal{K})$ whose intersections with $M - N(\mathcal{K})$ are pairwise disjoint. Then we form the surface H/\mathcal{D} by simultaneously surgering along each.

Lemma 2.9. *The previous definition is well defined, up to isotopy.*

Proof. If \mathcal{D} meets H in a collection of pairwise non-parallel curves, then this collection is unique, and thus H/\mathcal{D} is unique up to isotopy. Suppose, then, that two disks, D and E , in \mathcal{D} meet H in parallel curves. Then D and E can potentially be isotoped to meet H in two different arrangements. Surgery of H along each such arrangement produces surfaces H_1 and H_2 that differ only by a sphere. By the assumed irreducibility of M , this sphere bounds a ball. To form H/DE we remove this sphere from either H_1 or H_2 . In either case we get the same surface (up to isotopy), and thus the operation is well-defined.

The only other case that needs to be considered is when D and E are both ∂ -compressions. In this case the surfaces H_1 and H_2 described above differ by a disk with boundary on ∂M . By the assumed incompressibility of ∂M , and the irreducibility of M , this disk lies in a ball in M . It is thus removed from either H_1 or H_2 to form H/DE . In either case, the surface thus obtained is the same, up to isotopy. \square

Suppose now H is a properly embedded surface in M , transverse to \mathcal{T}^1 , and \mathcal{E} is a collection of interior edge-compressing disks for (H, \mathcal{T}^1) . Then the surface H/\mathcal{E} is defined only to be a surface in $M - N(\mathcal{T}^1)$. We now give a related operation that produces a surface $H//\mathcal{E}$ which




FIGURE 2. The collection \mathcal{E} consists of two disks that meet at a point of \mathcal{T}^1 , but are disjoint in $M - N(\mathcal{T}^1)$. The surfaces H and $H//\mathcal{E}$ are isotopic in M .




FIGURE 3. The collection \mathcal{E} consists of four disks that all meet at a point of \mathcal{T}^1 . The surface $H//\mathcal{E}$ can be obtained from H by surgery along a single compressing disk.

is properly embedded in M and transverse to T^1 . The surfaces H and $H//\mathcal{E}$ may be isotopic in M , as in Figure 2, or $H//\mathcal{E}$ may be obtained from H by surgery along some number of compressing disks, as in Figure 3.

Definition 2.10. Suppose \mathcal{E} is a collection of interior edge-compressing disks for (H, \mathcal{T}^1) which are pairwise disjoint in the complement of $N(\mathcal{T}^1)$. If, for each interior edge e of \mathcal{T}^1 , the surface H/\mathcal{E} meets $\partial N(e)$ in essential loops, then we cap off all such loops with disks in $N(\mathcal{T}^1)$ to form a properly embedded surface $H//\mathcal{E}$ in M that is transverse to \mathcal{T}^1 .

We define an equivalence relation on disks as follows:

Definition 2.11. Disks $C, C' \in \mathcal{D}(H, \mathcal{K})$ are *equivalent*, $C \sim C'$, if C and C' are isotopic in $M - N(\mathcal{K})$, through disks in $\mathcal{D}(H, \mathcal{K})$.

Definition 2.12. The *disk complex* $[\mathcal{D}(H, \mathcal{K})]$ is the complex defined as follows: vertices correspond to equivalence classes $[D]$ of the disk set $\mathcal{D}(H, \mathcal{K})$. A collection of n vertices spans an $(n - 1)$ -simplex if there are representatives of each whose intersections with $M - N(\mathcal{K})$ are pairwise disjoint.

Definition 2.13. Suppose τ is a simplex of $[\mathcal{D}(H, \mathcal{K})]$. Then $H/\tau = H/\mathcal{D}$, where \mathcal{D} is a pairwise disjoint collection of representatives of the vertices of τ .

Definition 2.14. The *homotopy index* of a complex $[\mathcal{D}(H, \mathcal{K})]$, denoted $\text{ind}[\mathcal{D}(H, \mathcal{K})]$, is defined to be 0 if $[\mathcal{D}(H, \mathcal{K})] = \emptyset$, and the smallest n such that $\pi_{n-1}([\mathcal{D}(H, \mathcal{K})])$ is non-trivial, otherwise. (If a complex is contractible, its homotopy index is left undefined.)

Definition 2.15. [Bac10] A surface H in a 3-manifold is *topologically minimal* if $[\mathcal{C}(H)]$ is either empty or non-contractible. When H is topologically minimal, we say the *index of H* is the number $\text{ind}[\mathcal{C}(H)]$.

We now come to our main theorem, which guarantees a process by which we can begin with an index n , topologically minimal surface H_0 , and end up with a surface H where $\text{ind}[\mathcal{CE}(H, \mathcal{T}^2)] \leq n$.

Theorem 2.16. *Suppose H_0 is a topologically minimal, non-peripheral surface in M whose index is n . Then there is a surface H , obtained from H_0 by the sequence of ∂ -compressions, interior edge-compressions, and isotopies given in the flow chart of Figure 4, such that*

$$\text{ind}[\mathcal{CE}(H, \mathcal{T}^2)] \leq n.$$

The proof of Theorem 1.2 is then complete by applying the following theorem.

Theorem 8.1. *Suppose $\text{ind}[\mathcal{CE}(H, \mathcal{T}^2)] = n$. Then*

$$\sum_{\Delta \in \mathcal{T}^3} \text{ind}[\mathcal{CE}(H \cap \Delta, \mathcal{T}^2)] = n.$$

Some remarks are in order about the flow chart of Figure 4.

Remark 2.17. The first loop of the flow chart concerns the set of ∂ -compressions, and thus only applies to surfaces with non-empty boundary. Theorem 3.3 in this loop always assumes the surface H that it is given has essential boundary on M , and produces a surface H/\mathcal{D} that

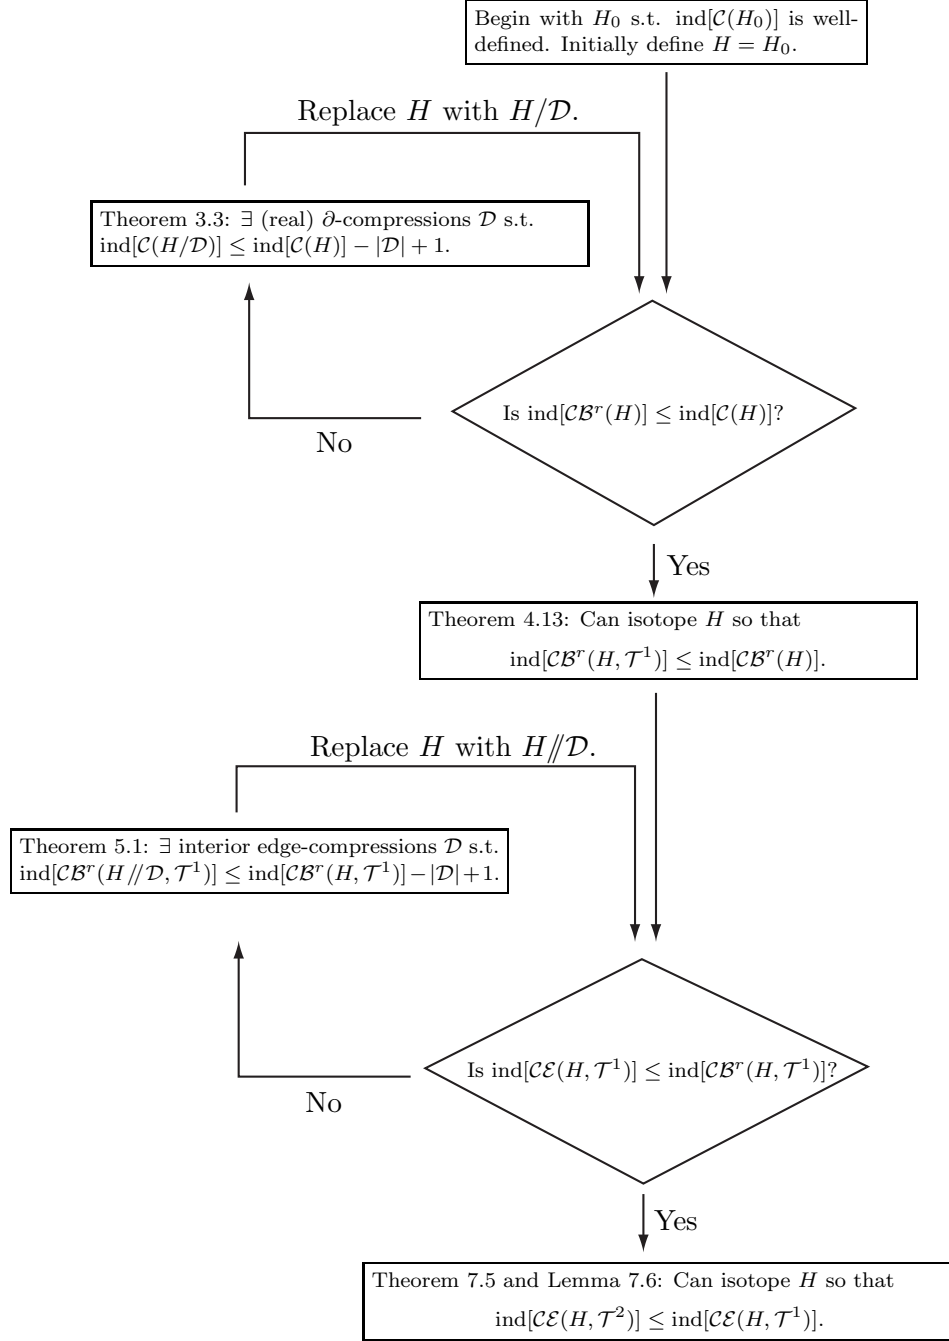


FIGURE 4. Transitioning from $[\mathcal{C}(H_0)]$ to $[\mathcal{CE}(H, \mathcal{T}^2)]$ in Theorem 2.16.

has essential boundary. When the genus of $\partial M \geq 2$, this allows us to bound the distance (as measured in the curve complex) between the boundary curves of the initial surface H_0 and the boundary curves of the surface that is eventually passed on to Theorem 4.13 in terms of the Euler characteristic of H_0 .

Remark 2.18. When $\partial M \cong \mathbb{T}^2$ the size of the collection \mathcal{D} given by Theorem 3.3 must be even. This is because each loop of both ∂H and $\partial(H/\mathcal{D})$ must be essential on ∂M , and thus the Euler characteristics of the components of $\partial M - H$ and $\partial M - (H/\mathcal{D})$ are all zero. Since surgery along a ∂ -compressing disk on one side of H decreases the Euler characteristic of the components of $\partial M - H$ on that side by precisely one and increases the Euler characteristic on the other side by one, there must be the same number of disks in \mathcal{D} on each side of H . Hence, $|\mathcal{D}|$ is even.

In this case one would measure the distance between the curves ∂H and $\partial(H/\mathcal{D})$ in the Farey graph. Since $|\mathcal{D}|$ is even it must be at least 2, and thus it follows from Theorem 3.3 that the distance between the boundary curves of the initial surface H_0 and the boundary curves of the surface that is eventually passed on to Theorem 4.13 is bounded by $\text{ind}[\mathcal{C}(H_0)]$. In particular, when $\text{ind}[\mathcal{C}(H_0)] = 2$, then these two sets of curves are at most distance 2 apart in the Farey graph. We make extensive use of this fact in [BDTSb].

Remark 2.19. In the second loop of the flow-chart, the collection \mathcal{D} produced by Theorem 5.1 has the property that H/\mathcal{D} meets the neighborhood of each interior edge of \mathcal{T}^1 in an essential loop, and thus $H//\mathcal{D}$ is well-defined. As in the previous remark, this implies that $|\mathcal{D}|$ is even. When $|\mathcal{D}| = 2$ the surface $H//\mathcal{D}$ is by construction isotopic to H in M . In particular, if the original surface H_0 is closed, and $\text{ind}[\mathcal{C}(H_0)] \leq 2$, then the final surface H produced by the flow chart will be isotopic to H_0 . This is why Theorem 1.1 is a special case of Theorem 1.2.

Remark 2.20. The proofs of all but one of the results indicated in the flow chart have appeared with different notation in previous work. In each case below we both indicate the original source, and reproduce the proof, adapted to the notation here. The one exception is Theorem 4.13, which is completely new. We consider this result to be the main technical heart of this paper. It plays a particularly important role here, as it links the two loops of the flow chart.

3. STAGE 1: $[\mathcal{C}(H)] \rightarrow [\mathcal{CB}^r(H)]$

Lemma 3.1. [cf. [Bacd], Theorem 4.5.] *Let X be a subcomplex of a simplicial complex Y . Suppose the homotopy index of X is n . Then either the homotopy index of Y is at most n , or there is a simplex $\tau \subset Y \setminus X$ such that the homotopy index of the subcomplex spanned by*

$$\{x \in X \mid \forall y \in \tau, x \text{ is adjacent to } y\}$$

is at most $n - \dim(\tau)$.

Proof. For every simplex τ of Y , let

$$V_\tau = \{x \in X \mid \forall y \in (\tau \setminus X), x \text{ is adjacent to } y\}.$$

Note that if $\tau \subset X$ then $\tau \setminus X = \emptyset$, and thus V_τ will be all of X .

For all $\tau \subset Y$, let $\tau^+ = \tau \setminus X$. Our goal is to show that there is a simplex $\tau \subset Y$ with $\tau^+ \neq \emptyset$ such that the homotopy index of $V_\tau = V_{\tau^+}$ is at most $n - \dim(\tau^+)$. If $X = Y$ then the result is immediate. If not, then there exists a simplex τ such that $\tau^+ \neq \emptyset$. By way of contradiction, we suppose that for all $\tau \subset Y$, V_τ does not have homotopy index less than or equal to $n - \dim(\tau^+)$. Thus, for all τ , $V_\tau \neq \emptyset$ and

$$(1) \quad \pi_i(V_\tau) = 1 \text{ for all } i \leq n - \dim(\tau^+) - 1.$$

Claim 3.2. *Suppose τ is a cell of Y which lies on the boundary of a cell σ . Then $V_\sigma \subset V_\tau$.*

Proof. Suppose $v \in V_\sigma$. Then $v \in X$ and v is connected by an edge to every vertex of $\sigma \setminus X$. Since τ lies on the boundary of σ , it follows that $\tau \setminus X \subset \sigma \setminus X$, and thus v is connected by an edge to every vertex of $\tau \setminus X$. Thus, $v \in V_\tau$. \square

If $X = \emptyset$ then the result is immediate, as V_τ will be \emptyset (and thus have homotopy index 0) for any simplex τ in Y .

If $X \neq \emptyset$ then, by assumption, there is a non-trivial map ι from an $(n-1)$ -sphere S into the $(n-1)$ -skeleton of X . Assuming the theorem is false will allow us to inductively construct a map Ψ of a n -ball B into X such that $\Psi(\partial B) = \iota(S)$. The existence of such a map contradicts the non-triviality of ι .

Since $X \subset Y$, ι is also a map from S into Y . If $\pi_{n-1}(Y) \neq 1$ then the result is immediate. Otherwise, ι can be extended to a map from an n -ball B into Y . Let Σ denote a triangulation of B so that the map $\iota : B \rightarrow Y$ is simplicial.

Push the triangulation Σ into the interior of B , so that $\text{Nbhd}(\partial B)$ is no longer triangulated (Figure 5(b)). Then extend Σ to a cell decomposition over all of B by forming the product of each cell of $\Sigma \cap S$ with the interval I (Figure 5(c)). Denote this cell decomposition as Σ' .




FIGURE 5. (a) The triangulation Σ of B . (b) Push Σ into the interior of B . (c) Fill in $\text{Nbhd}(\partial B)$ with product cells to complete Σ' . (d) Σ^* is the dual of Σ' .

Note that ι extends naturally over Σ' , and the conclusion of Claim 3.2 holds for cells of Σ' . Now let Σ^* denote the dual cell decomposition of Σ' (Figure 5(d)). This is done in the usual way, so that there is a correspondence between the d -cells of Σ^* and the $(n - d)$ -cells of Σ' . Note that, as in the figure, Σ^* is not a cell decomposition of all of B , but rather a slightly smaller n -ball, which we call B' .

For each cell τ of Σ' , let τ^* denote its dual in Σ^* , and let $V_\tau = V_{\iota(\tau)}$. Thus, it follows from Claim 3.2 that if σ^* is a cell of Σ^* that is on the boundary of τ^* , then $V_\sigma \subset V_\tau$.

We now produce a contradiction by defining a continuous map $\Psi : B' \rightarrow X$ such that $\Psi(\partial B') = \iota(S)$. The map will be defined inductively on the d -skeleton of Σ^* so that the image of every cell τ^* is contained in V_τ .

For each 0-cell $\tau^* \in \Sigma^*$, choose a point in V_τ to be $\Psi(\tau^*)$. If τ^* is in the interior of B' then this point may be chosen arbitrarily in V_τ . If $\tau^* \in \partial B'$ then τ is an n -cell of Σ' . This n -cell is $\sigma \times I$, for some $(n - 1)$ -cell σ of $\Sigma \cap S$. But since $\iota(S) \subset X$, it follows that $\iota(\tau) \subset X$, and thus $V_\tau = X$. We conclude $\iota(\tau) \subset V_\tau$, and thus we can choose $\iota(\tau^*)$, the barycenter of $\iota(\tau)$, to be $\Psi(\tau^*)$.

We now proceed to define the rest of the map Ψ by induction. Let τ^* be a d -dimensional cell of Σ^* and assume Ψ has been defined on the $(d - 1)$ -skeleton of Σ^* . In particular, Ψ has been defined on $\partial\tau^*$. Suppose σ^* is a cell on $\partial\tau^*$. By Claim 3.2 $V_\sigma \subset V_\tau$. By assumption $\Psi|_{\sigma^*}$ is defined and $\Psi(\sigma^*) \subset V_\sigma$. We conclude $\Psi(\sigma^*) \subset V_\tau$ for all $\sigma^* \subset \partial\tau^*$, and thus

$$(2) \quad \Psi(\partial\tau^*) \subset V_\tau.$$

Note that

$$\dim(\tau) = n - \dim(\tau^*) = n - d.$$

Since $\dim(\tau^+) \leq \dim(\tau)$, we have

$$\dim(\tau^+) \leq n - d.$$

Thus

$$d \leq n - \dim(\tau^+),$$

and finally

$$d - 1 \leq n - \dim(\tau^+) - 1.$$

It now follows from Equation 1 that $\pi_{(d-1)}(V_\tau) = 1$. Since $d - 1$ is the dimension of $\partial\tau^*$, we can thus extend Ψ to a map from τ^* into V_τ .

Finally, we claim that if $\tau^* \subset \partial B'$ then this extension of Ψ over τ^* can be done in such a way so that $\Psi(\tau^*) = \iota(\tau^*)$. This is because in this case each vertex of $\iota(\tau)$ is in X , and hence $V_\tau = X$. As $\iota(S) \subset X = V_\tau$, we have $\iota(\tau^*) \subset V_\tau$. Thus we may choose $\Psi(\tau^*)$ to be $\iota(\tau^*)$. \square

Theorem 3.3. *Let H be a non-peripheral, topologically minimal surface in M . Then either $\text{ind}[\mathcal{CB}^r(H)] \leq \text{ind}[\mathcal{C}(H)]$, or there exists a collection \mathcal{D} of pairwise disjoint, real ∂ -compressing disks for H such that*

$$\text{ind}[\mathcal{C}(H/\mathcal{D})] \leq \text{ind}[\mathcal{C}(H)] - |\mathcal{D}| + 1.$$

Proof. Set $X = [\mathcal{C}(H)]$, $Y = [\mathcal{CB}^r(H)]$, and apply Lemma 3.1. The conclusion of the lemma gives us two possibilities. The first is that $\text{ind}[\mathcal{CB}^r(H)] \leq \text{ind}[\mathcal{C}(H)]$. The second is that there is a simplex $\tau \subset [\mathcal{B}^r(H)]$ such that the homotopy index of the complex Z spanned by

$$\{x \in [\mathcal{C}(H)] \mid \forall y \in \tau, x \text{ is adjacent to } y\}$$

is at most $\text{ind}[\mathcal{C}(H_m)] - \dim(\tau)$. Let \mathcal{D} denote a pairwise disjoint collection of disks representing the vertices of τ . Then $\dim(\tau) = |\mathcal{D}| - 1$. We claim that Z is precisely $[\mathcal{C}(H/\mathcal{D})]$, and thus we have obtained the desired conclusion.

Note that the vertices of Z correspond to the set of compressing disks for H that are disjoint from every disk in \mathcal{D} . Let H' denote the surface obtained from H by simultaneous surgery along the disks in \mathcal{D} , so that H/\mathcal{D} is obtained from H' by removing any component that lies in a ball. (If $H/\mathcal{D} = \emptyset$ then the surface H would have been peripheral, a contradiction.) It follows immediately that $\mathcal{C}(H') = Z$. What remains is to show $\mathcal{C}(H') = \mathcal{C}(H/\mathcal{D})$.

We now claim that every boundary component of every non-disk component of H' is essential on ∂M . If not, then such a boundary component that is innermost on ∂M bounds a subdisk C of ∂M that is isotopic to a compressing disk for H' . But, as this disk is isotopic into ∂M , it can be made disjoint from every other disk in $\mathcal{C}(H')$. Thus, $[\mathcal{C}(H')]$ would be a contractible complex, a contradiction.

Suppose now E is a compressing disk for H' . Then E meets a component H'' of H' that is not a disk. The surface H'' will either be closed, and hence had been unaffected by the surgery along the disks in \mathcal{D} , or by the above argument will have boundary that is essential on ∂M . As ∂M is assumed to be incompressible, we conclude $\partial H''$ consists of essential loops in M . Hence, H'' can not be contained in a ball, and is thus a component of H/\mathcal{D} . We conclude E is a compressing disk for H/\mathcal{D} , and thus $\mathcal{C}(H') = \mathcal{C}(H/\mathcal{D})$. \square

4. STAGE 2: $[\mathcal{CB}^r(H)] \rightarrow [\mathcal{CB}^r(H, \mathcal{T}^1)]$

Theorem 3.3 has left us with a surface H such that $[\mathcal{CB}^r(H)]$ has well-defined homotopy index. The goal of this section is to show that H can be isotoped so that $\text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)] \leq \text{ind}[\mathcal{CB}^r(H)]$.

Henceforth, we will assume that H has been isotoped so that $|\partial H \cap \mathcal{T}^1|$ is minimal. This has the following nice consequence.

Lemma 4.1. *If $|\partial H \cap \mathcal{T}^1|$ is minimal and $B \in \mathcal{B}^r(H, \mathcal{T}^1)$, then $B \in \mathcal{B}^r(H)$.*

Proof. If $B \notin \mathcal{B}^r(H)$ then either B is not a ∂ -compressing disk for H , or B is fake. In the former case B cuts off a subdisk C of H . As ∂M is incompressible, and $C \cup B$ is a disk incident to ∂M , it follows that B meets ∂M in an arc β that is parallel to an arc α of ∂H . By definition, this is also true when B is fake. Since $B \in \mathcal{B}^r(H, \mathcal{T}^1)$, β does not meet \mathcal{T}^1 . Hence, isotoping α to β reduces $|\partial H \cap \mathcal{T}^1|$, a contradiction. \square

Definition 4.2. Suppose $D \in \mathcal{CB}^r(H, \mathcal{T}^1)$. If $D \in \mathcal{CB}^r(H)$ then we say it is *honest*. Otherwise, we say D is *dishonest*.

By Lemma 4.1, if $|\partial H \cap \mathcal{T}^1|$ is minimal and $D \in \mathcal{B}^r(H, \mathcal{T}^1)$ then D must be honest. It follows that if D is a dishonest element of $\mathcal{CB}^r(H, \mathcal{T}^1)$ then $D \in \mathcal{C}(H, \mathcal{T}^1)$ but $D \notin \mathcal{C}(H)$. Thus, ∂D bounds a subdisk C of H whose interior meets \mathcal{T}^1 .

Lemma 4.3. *Suppose D and E are disjoint elements of $\mathcal{CB}^r(H, \mathcal{T}^1)$ and D is dishonest. Then $E \in \mathcal{CB}^r(H/D, \mathcal{T}^1)$.*

Proof. H/D is obtained from H by removing a neighborhood of $D \cap H$, and replacing it with two copies D' and D'' of D . If the lemma is false, then $E \cap (H/D)$ cuts off a subdisk E^* of H/D that is disjoint from \mathcal{T}^1 . If E^* does not contain D' or D'' , then E would not have been a compression for $H - N(\mathcal{T}^1)$. If E^* contains both D' and D'' , then D would have been honest. Finally, if E^* contains one of D' or D'' (assume the former), then $D \cap H$ is isotopic to $E \cap H$ on $H - N(\mathcal{T}^1)$. When we isotope these intersections to coincide, then $D \cup E$ becomes

either a sphere in $M - N(\mathcal{T}^1)$, or a disk with boundary on $\partial M - N(\mathcal{T}^1)$. In either case, $D \cup E$ cuts a ball B out of M . Since D is assumed to be dishonest, it follows that E is dishonest as well. Thus, the ball B must contain all of \mathcal{T}^1 , a contradiction. \square

We now define a complexity on H , and show that both honest and dishonest compressions and ∂ -compressions decrease complexity.

Definition 4.4. If H is empty, then we define the *width* $w(H)$ to be $(0, 0)$. If H is connected, then the *width* $w(H)$ is the pair $(-\chi(H), |\mathcal{T}^1 \cap H|)$. If H is disconnected, then its *width* is the ordered set of the widths of its components, where we include repeated pairs and the ordering is non-increasing. Comparisons are made lexicographically at all levels.

Lemma 4.5. *Suppose $D \in \mathcal{CB}^r(H, \mathcal{T}^1)$. Then $w(H/D) < w(H)$.*

Proof. Suppose first that D is honest, so that $D \cap H$ is essential in H . If, furthermore, ∂D does not separate H , then $-\chi(H/D)$ is less than $-\chi(H)$, and hence the width is also less. If, on the other hand, ∂D separates H then H/D is disconnected, and both components have smaller negative Euler characteristic. Hence, again width has decreased.

Now suppose D is dishonest. Then ∂D cuts off a subdisk D' of H whose interior meets \mathcal{T}^1 . Thus, $D' \cap \mathcal{T}^1 \neq \emptyset$. The result of surgering along D produces a sphere and a surface that has the same Euler characteristic as H , but meets \mathcal{T}^1 fewer times, and thus has smaller width than H . This latter surface is precisely H/D . \square

Definition 4.6. For any simplicial complex X , let $X^{(1)}$ denote the first barycentric subdivision of X . If v is a vertex of $X^{(1)}$, then v^* will denote the simplex of X that v is the barycenter of.

Definition 4.7. A *compressing n -sequence* Σ for a surface H is a simplicial n -complex (homeomorphic to D^2 when $n = 2$) where each vertex v is associated with a surface H_v in M which is either isotopic to H in M , or is obtained from H by compression or ∂ -compression. Let Σ_0 denote the subcomplex of Σ spanned by those vertices v for which H_v is isotopic to H in M . Σ admits compatible global and local structures, as follows.

Global Structure

- (1) When $v \in \Sigma_0$ there is a fixed isotopy from H_v to H , inducing an identification of $[\mathcal{CB}^r(H_v)]$ with $[\mathcal{CB}^r(H)]$.
- (2) There is a simplicial map $\iota : \Sigma - \Sigma_0 \rightarrow [\mathcal{CB}^r(H)]^{(1)}$, such that for each v in the domain of ι the surface H_v is isotopic in M to $H/\iota(v)^*$. See Figure 6.




FIGURE 6. The unshaded region of Σ represents Σ_0 . The lighter shaded region is $\Sigma - \Sigma_0$, the domain of ι .

Local Structure

- (1) If e is an edge of Σ with endpoints u and v , where $H_u, H_v \in \Sigma_0$ and $w(H_u) = w(H_v)$, then $H_u = H_v$.
- (2) Let v be a vertex in Σ_0 , and $\text{link}(v)^-$ the subcomplex of $\text{link}(v)$ spanned by vertices u such that $w(H_u) < w(H_v)$. Then there is a simplicial map $\eta_v : \text{link}(v)^- \rightarrow [\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$ such that for each $w \in \text{link}(v)^-$, $H_w = H_v / \eta_v(w)^*$.¹ See Figure 7.

Compatibility Condition

For each vertex v in Σ_0 , and each $w \in \text{link}(v)^-$, the simplex of $[\mathcal{CB}^r(H)]$ spanned by the vertices of $\eta_v(w)^*$ that represent honest compressing disks for H is precisely the simplex $\iota(w)^*$. See Figure 8.

¹Equality here indicates that there exists a set of disk representatives for the vertices of $\eta_v(w)^*$ so that H_w is precisely the surface obtained from H_v by surgering along these disks.




FIGURE 7. $\text{link}(v)^-$ is the subcomplex of $\text{link}(v)$ spanned by vertices u for which $w(H_u) < w(H_v)$. The map η_v sends $\text{link}(v)^-$ into $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$.

Definition 4.8. Suppose Σ is a compressing n -sequence, $v \in \Sigma_0$, and V is a component of the subcomplex of Σ such that $H_w = H_v$ for all $w \in V$. Let $\text{star}(V)$ denote the union of the simplices of Σ that meet V , $\overline{\text{star}}(V)$ the closure of $\text{star}(V)$, and $\text{link}(V) = \overline{\text{star}}(V) - \text{star}(V)$. Then we say V is a *plateau* if $w(H_v) > w(H_u)$ for all $u \in \text{link}(V)$.

Now suppose V is a plateau, and v is a vertex of V . Then note that the maps η_w , for each $w \in V$, induce a simplicial map $\eta_v : \text{link}(V) \rightarrow [\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$. See Figure 9.

Definition 4.9. Let S be an $(n - 1)$ -cycle in the singular homology of $[\mathcal{CB}^r(H)]^{(1)}$. We say a compressing n -sequence Σ for H *spans* S if $\iota(\partial\Sigma) = S$.




FIGURE 8. The compatibility condition implies that $\iota(w)$ is at the barycenter of the simplex of $[\mathcal{CB}^r(H)]^{(1)}$ spanned by the vertices of $\eta_v(w)^*$ corresponding to honest compressing disks.

Lemma 4.10. *Suppose Σ is a compressing n -sequence that spans an $(n-1)$ -cycle S . If $n = 2$ and S is non-trivial in $\pi_1([\mathcal{CB}^r(H)]^{(1)})$, or $n \neq 2$ and S is non-trivial in $H_{n-1}([\mathcal{CB}^r(H)]^{(1)})$, then Σ has a plateau V .*

Proof. First, note that Σ_0 can not be empty because otherwise ι would map all of Σ into $[\mathcal{CB}^r(H)]^{(1)}$. When $n = 2$ this gives a homotopy of S to a point, and when $n \neq 2$ this gives a singular n -chain that S is the boundary of. Now let v denote a vertex of Σ_0 , for which $w(H_v)$ is maximal. Let V denote the component of the subcomplex of Σ spanned by vertices w with $H_w = H_v$, containing v . Then V must be a plateau. \square




FIGURE 9. The darker shaded region represents a plateau V . The map η_V sends $\text{link}(V)$ into $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$, where v is any vertex of V .

Definition 4.11. Let Σ be a compressing n -sequence. Let $\{v_i\}$ denote a choice of vertex in each plateau. Then the *size* of Σ is the ordered set $\{w(H_{v_i})\}$, where the ordering is non-increasing and comparisons are made lexicographically.

By Lemma 4.10, the size of a compressing n -sequence is well-defined, as long as it spans a homotopically/homologically non-trivial cycle S .

Lemma 4.12. *Let S be an $(n-1)$ -cycle, and Σ a minimal size compressing n -sequence that spans S . Let V be a plateau of Σ . If $n = 2$ and S is non-trivial in $\pi_1([\mathcal{CB}^r(H)]^{(1)})$ then $\eta_V(\text{link}(V))$ is non-trivial in $\pi_1([\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)})$. If $n \neq 2$ and S is non-trivial in $H_{n-1}([\mathcal{CB}^r(H)]^{(1)})$, then $\eta_V(\text{link}(V))$ is non-trivial in $H_{n-1}([\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)})$.*

Proof. We first address the case where $n = 2$. Let α be a component of $\text{link}(V)$ that is outermost on the disk Σ . Let A be the subdisk of Σ bounded by α , and note that $V \subset A$. If the lemma is false, then $\eta_V(\alpha)$ is trivial in $\pi_1([\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)})$. Thus, we may extend the map $\eta_V|_\alpha$ over a disk Π . That is, there is a map $\phi : \Pi \rightarrow \pi_1([\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)})$ such that $\phi(\partial\Pi) = \eta_V(\alpha)$. In this case, we show we can create a new compressing n -sequence of smaller size by replacing the disk A in Σ with Π .

If $n \neq 2$ and the lemma is false, then η_V can be extended to a simplicial map ϕ from a simplicial n -complex Π into $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$

(so that $\phi(\partial\Pi) = \eta_V(\text{link}(V))$). In this case, we show we can create a new compressing n -sequence of smaller size by replacing $\overline{\text{star}}(V)$ in Σ with Π .

In either case, to realize the new complex thus obtained as a compressing n -sequence, we go through each of the conditions in Definition 4.7. Choose $v \in V$. We associate each vertex w of Π with a surface H_w in M by letting $H_w = H_v/\phi(w)^*$. Note that when H_w is isotopic to H_v (and thus H), this gives a particular isotopy.

To satisfy the global condition of Definition 4.7, we must extend the map ι over those vertices w of Π that are now associated with surfaces not isotopic to H . We will call the new map thus created ι' . This map is defined by letting $\iota'(w)$ be the barycenter of the simplex spanned by the vertices of $\phi(w)^*$ representing honest compressing disks. Since $\phi(\partial\Pi) = \eta_V(\partial\Pi)$, the compatibility condition for η_V and the original map ι guarantees that $\iota'|\partial\Pi = \iota|\partial\Pi$.

We now discuss the local condition. Let w be a vertex of Π for which H_w is isotopic to H in M . Let u be a vertex adjacent to w . Since u and w are connected by an edge, and the map ϕ is simplicial, we may conclude that $\phi(u)$ and $\phi(w)$ are either the same vertex of $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$, or they are connected by an edge in $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$. In the former case H_u and H_w will be the same surface. In the latter case one of $\phi(u)^*$ or $\phi(w)^*$ is a face of the other. If $\phi(u)^*$ is a face of $\phi(w)^*$ then by Lemma 4.5 the width of $H_u = H_v/\phi(u)^*$ is larger than the width of $H_w = H_v/\phi(w)^*$. In this case there is nothing further to show. When $\phi(w)^*$ is a face of $\phi(u)^*$ the width of H_u is less than the width of H_w , and hence $u \in \text{link}(w)^-$. In this case we must now define $\eta_w(u)$. Since H_w is isotopic to H in M , it must be the case that the disks represented by the vertices of $\phi(w)^*$ are dishonest. Hence, by Lemma 4.3 H_u is obtained from H_w by surgering along the simplex $\phi(u)^* - \phi(w)^*$. We may thus define $\eta_w(u)$ to be the barycenter of the simplex $\phi(u)^* - \phi(w)^*$, so that $H_u = H_w/\eta_w(u)^*$.

Finally, we must check the compatibility of the maps η_w and ι' defined above. Since all of the compressing disks represented by vertices of $\phi(w)^*$ are dishonest, the vertices of $\eta_w(u)^* = \phi(u)^* - \phi(w)^*$ that represent honest compressing disks are just the vertices of $\phi(u)^*$ that represent honest compressing disks. These vertices span the simplex that $\iota'(u)$ is the barycenter of, as required by the compatibility condition. \square

Theorem 4.13. *If $\text{ind}[\mathcal{CB}^r(H)]$ is well defined, then H may be isotoped so that*

$$[\mathcal{CB}^r(H, \mathcal{T}^1)] \leq \text{ind}[\mathcal{CB}^r(H)].$$




FIGURE 10. (a) A simplex σ of Φ' . (b) The simplex $\sigma * z$ of Φ . (c) Realizing $\sigma * z$ as a cube C of $\bar{\Phi}$.


Proof. By assumption there is a homotopically non-trivial map ι from an $(n-1)$ -sphere S into $[\mathcal{CB}^r(H)]$. When $n \neq 2$, it is a consequence of the Hurewicz Theorem that this map is homologically non-trivial. In any case, our goal is to construct a compressing n -cycle that spans S . It then follows from Lemmas 4.10 and 4.12 that a minimal such compressing n -sequence has a vertex v such that H_v is isotopic to H in M and the homotopy index of $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]$ is at most n . The result thus follows.

Let Φ' denote a triangulation of S so that the map ι is simplicial. Let B denote the n -ball obtained by coning S to a point z . By coning each simplex of Φ' to z we thus get a triangulation Φ of B .

We now turn the triangulation Φ into a cubing $\bar{\Phi}$ of B . Let σ be an $(n-1)$ -simplex of Φ' with vertices $\{v_i\}_{i=1}^n$. Let e_i denote an orthonormal basis for \mathbb{R}^n , and let C denote the unit cube spanned by these vectors. Then we can identify the simplex $\sigma * z$ of Φ with C by placing

- z at the origin,
- v_i at the corner of C given by the vector e_i , and
- the barycenter of each face σ' of σ at the corner of C given by $\sum e_i$, where the sum is taken over all i such that v_i is a vertex of σ' .

See Figure 10. Note that the map ι , which sends the simplices of $\Phi' = \Phi|S$ to simplices of $[\mathcal{CB}^r(H)]$, now extends naturally to a map that sends simplices of $\bar{\Phi}|S$ to simplices of $[\mathcal{CB}^r(H)]^{(1)}$. We continue to denote this map ι .

FIGURE 11. The surfaces H_i^j defined by a disk D_i .

We now subdivide the cubing $\bar{\Phi}$ to get a finer cubing of B , and associate a surface in M with each vertex of this finer cubing. To begin, choose a representative disk for each vertex of $\iota(S)$, so that if τ is a simplex in the image of ι , then the chosen representatives of the vertices of τ are pairwise disjoint. For the cube C above, let D_i denote the chosen representative of v_i . For each i , let $n_i = |\mathcal{T}^1 \cap D_i|$. In each disk D_i choose a collection of arcs $\{\alpha_i^j\}_{j=1}^{n_i}$ such that

- (1) For each i and j , α_i^j connects a point of $D_i \cap \mathcal{T}^1$ to a point in ∂D_i .
- (2) For each j and k where $j < k$, the arcs α_i^j and α_i^k should either be disjoint, or $\alpha_i^j \subset \alpha_i^k$.
- (3) If D_i meets ∂M , then there is a j such that $\alpha_i^j \subset (D_i \cap \partial M)$, and contains all of the points of $\partial D_i \cap \mathcal{T}^1$.

Finally, choose a small enough neighborhood $N(D_i)$ so that if $D_i \cap D_j = \emptyset$, then $N(D_i) \cap N(D_j) = \emptyset$, and let $H_i = H \cap N(D_i)$.

Let $H_i^0 = H_i$. For each $1 \leq J \leq n_i$, let H_i^J denote the surface obtained from H_i by simultaneous surgery along all of the components of $\bigcup_{j=1}^J \alpha_i^j$. Note that H_i/D_i can then be obtained from $H_i^{n_i}$ by an honest

compression or ∂ -compression, and for each $j \geq 1$ the surface H_i^{j-1} can be obtained from H_i^j by a dishonest compression or ∂ -compression. See Figure 11.

We now subdivide the cube C of $\bar{\Phi}$ into smaller subcubes. For each i , cut C by n_i planes orthogonal to \mathbf{e}_i . Each vertex of this new cubing is then of the form

$$x = \sum_{i=1}^n \frac{x(i)}{n_i + 1} \mathbf{e}_i$$




FIGURE 12. (a) The subcube c , when $\mathcal{D} = \{A, B, C\}$.
 (b) The simplex in c defined by the sequence $\mathcal{D}_0 = \emptyset$, $\mathcal{D}_1 = \{A\}$, $\mathcal{D}_2 = \{A, B\}$, and $\mathcal{D}_3 = \mathcal{D}$.

where $0 \leq x(i) \leq n_i + 1$ for each i . We associate to the vertex x the surface H_x in M obtained from H by replacing H_i with

- $H_i^{x(i)}$, if $x(i) \leq n_i$
- H_i/D_i , if $x(i) = n_i + 1$

for each i .

We now triangulate each of the subcubes of C . Suppose x and y are vertices of such a subcube c that are connected by an edge. Then $x(i)$ and $y(i)$ differ by exactly one for one value j of i , and are equal for every $i \neq j$. Hence, the surfaces H_x and H_y are identical away from $N(D_j)$. Since $x(j)$ and $y(j)$ differ by exactly one, it follows that one of $H_j^{x(j)}$ and $H_j^{y(j)}$ is obtained from the other by a compression in $N(D_j)$. (This compression will be honest if either $x(j)$ or $y(j)$ is $n_j + 1$, and dishonest otherwise.) Furthermore, parallel edges of the subcube c will also correspond to pairs of surfaces obtained by compressing along the same disk D_j . It follows that for the subcube c there is a vertex v , represented by a surface H_v , and a collection \mathcal{D} of honest and dishonest compressing disks for H_v , such that every other vertex is obtained from H_v by simultaneous compression along some subcollection of \mathcal{D} . See Figure 12(a).

We now break up each subcube c of C into a collection of n -simplices, thereby obtaining a triangulation of B . Let $\{\mathcal{D}_i\}_{i=0}^n$ denote a sequence of subsets of \mathcal{D} , such that

- (1) $\mathcal{D}_n = \mathcal{D}$
- (2) \mathcal{D}_i is obtained from \mathcal{D}_{i+1} by removing one disk.

Thus $\mathcal{D}_0 = \emptyset$ and for each $i \geq 1$, \mathcal{D}_i is a collection with precisely i disks. Any such sequence $\{\mathcal{D}_i\}$ defines a simplex in c spanned by the vertices of c associated to the surfaces H_v/\mathcal{D}_i . See Figure 12(b).

To show that the simplicial complex thus defined, with the given associated surfaces for each vertex, is a compressing n -sequence spanning S we make a few observations. First, note that by construction, every vertex besides those in ∂B are isotopic in M to H . Thus, Σ_0 is everything besides $\partial B = S$. The original map ι which sent S into $[\mathcal{CB}^r(H)]$ has been refined by our construction to be a map from $S^{(1)}$ into $[\mathcal{CB}^r(H)]^{(1)}$, and this is precisely the map required by the global condition of Definition 4.7.

Now consider a vertex v of Σ_0 . The surface H_v is associated to a surface obtained from H by surgering along some collection of arcs. For each vertex w in $\text{link}(v)^-$, the surface H_w is obtained from H by surgering along a subcollection of these arcs, and possibly along some honest compressing disks for H_v . In other words, H_w is obtained from H_v by a collection of honest and dishonest compressing disks. We may thus define $\eta_v(w)$ to be the barycenter of the simplex of $[\mathcal{CB}^r(H_v, \mathcal{T}^1)]^{(1)}$ spanned by these compressing disks. Compatibility with the map ι then follows by construction. \square

5. STAGE 3: $[\mathcal{CB}^r(H, \mathcal{T}^1)] \rightarrow [\mathcal{CE}(H, \mathcal{T}^1)]$

Theorem 5.1. *Let H be a surface in M such that $\text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)]$ is well defined. Then either $\text{ind}[\mathcal{CE}(H, \mathcal{T}^1)] \leq \text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)]$, or there is a collection \mathcal{D} of interior edge-compressing disks for H such that*

$$\text{ind}[\mathcal{CB}^r(H//\mathcal{D}, \mathcal{T}^1)] \leq \text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)] - |\mathcal{D}| + 1.$$

Proof. Since ∂H is normal, every element of $\mathcal{B}^r(H, \mathcal{T}^1)$ is isotopic to at least one element of $\mathcal{E}(H, \mathcal{T}^1)$. Let X denote the subcomplex of $\mathcal{E}(H, \mathcal{T}^1)$ spanned by those disks that are isotopic to elements of $\mathcal{B}^r(H, \mathcal{T}^1)$, together with the complex $[\mathcal{C}(H, \mathcal{T}^1)]$. Then X is homotopy equivalent (in fact, homeomorphic) to $[\mathcal{CB}^r(H, \mathcal{T}^1)]$. If we set $Y = [\mathcal{CE}(H, \mathcal{T}^1)]$, then $X \subset Y$. We may thus apply Lemma 3.1. The conclusion of this lemma gives us two possibilities. The first is that $\text{ind}[\mathcal{CE}(H, \mathcal{T}^1)] \leq \text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)]$, and thus we immediately obtain the desired result. The second possibility is that there is a simplex τ

spanned by edge-compressing disks incident to interior edges, such that the homotopy index of the complex Z spanned by

$$\{x \in [\mathcal{CB}^r(H, \mathcal{T}^1)] \mid \forall y \in \tau, x \text{ is adjacent to } y\}$$

is at most $\text{ind}[\mathcal{CB}^r(H, \mathcal{T}^1)] - \dim(\tau)$.

Let \mathcal{D} denote a collection of disks representing the vertices of τ which is pairwise disjoint in the complement of $N(\mathcal{T}^1)$. Then $\dim(\tau) = |\mathcal{D}| - 1$. Let H' denote the surface in $M - N(\mathcal{T}^1)$ obtained from $H - N(\mathcal{T}^1)$ by simultaneous surgery along the disks in \mathcal{D} . Each element of Z can then be identified with a compressing or real ∂ -compressing disk for H' . We claim that $H//\mathcal{D}$ is well-defined, and thus $Z = [\mathcal{CB}^r(H//\mathcal{D}, \mathcal{T}^1)]$. The desired result follows.

To show $H//\mathcal{D}$ is well-defined, we must prove that for every interior edge e of \mathcal{T}^1 , every component of $H' \cap \partial N(e)$ is essential. If not, then such a boundary component that is innermost on $\partial N(e)$ bounds a subdisk C of $\partial N(e)$ that is isotopic to a compressing disk for H' . But, as this disk is isotopic into $\partial N(\mathcal{T}^1)$, it can be made disjoint from every other disk in Z . Thus, Z would be a contractible complex, a contradiction. \square

6. STAGE 4: $[\mathcal{CE}(H, \mathcal{T}^1)] \rightarrow [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$

Definition 6.1. Suppose $E \in \mathcal{CE}(H, \mathcal{T}^1)$. We say E' is a *shadow* of E if $\partial E' = \partial E$, and E' is disjoint from \mathcal{T}^2 away from its boundary. (The interior of a shadow E' may intersect H .) We define the set $\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)$ to be the subset of $\mathcal{CE}(H, \mathcal{T}^1)$ consisting of those disks that have shadows.

Theorem 6.2. [cf. [Bac10], Theorem 3.2.] *Suppose the homotopy index of $[\mathcal{CE}(H, \mathcal{T}^1)]$ is n . Then H may be isotoped (rel \mathcal{T}^1) so that*

- (1) *H meets the 2-simplices of \mathcal{T}^2 in p points of tangency, for some $p \leq n$. Away from these tangencies H is transverse to \mathcal{T}^2 .*
- (2) *The complex $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ has homotopy index $i \leq n - p$.*

Proof. When $\text{ind}[\mathcal{CE}(H, \mathcal{T}^1)] = 0$ the result is immediate, as $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)] \subset [\mathcal{CE}(H, \mathcal{T}^1)]$. We will assume, then, that $\text{ind}[\mathcal{CE}(H, \mathcal{T}^1)] = n \geq 1$. It follows that $\pi_{n-1}([\mathcal{CE}(H, \mathcal{T}^1)]) \neq 1$, and thus there is a map $\iota : S \rightarrow [\mathcal{CE}(H, \mathcal{T}^1)]$ of an $(n-1)$ -sphere S into the $(n-1)$ -skeleton of $[\mathcal{CE}(H, \mathcal{T}^1)]$ which is not homotopic to a point. Let B be the cone on S to a point z . (The point z is necessarily not in $[\mathcal{CE}(H, \mathcal{T}^1)]$.) Hence, B is an n -ball.

Our first challenge is to define a continuous family of surfaces H_x in M isotopic to H , where $x \in B$. Let Σ be a triangulation of $S = \partial B$ so that the map ι is simplicial. Let $\{v_i\}$ denote the set of vertices of

$[\mathcal{CE}(H, \mathcal{T}^1)]$ that are contained in $\iota(S)$. For each i choose a representative D_i from the equivalence class of disks represented by v_i so that if (v_i, v_j) is an edge of $[\mathcal{CE}(H, \mathcal{T}^1)]$, then $D_i \cap D_j = \emptyset$.

For each i , let N_i be a neighborhood of D_i in $M - N(\mathcal{T}^1)$. Assume N_i has been chosen to be small enough so that $N_i \cap N_j = \emptyset$ whenever (v_i, v_j) is an edge of $[\mathcal{CE}(H, \mathcal{T}^1)]$. Let $H_i(t)$ be a family of surfaces in N_i so that $H_i(0) = H \cap N_i$, and $H_i(1)$ is obtained from $H \cap N_i$ by shrinking the disk D_i until it meets \mathcal{T}^2 in at most a subarc of an edge of \mathcal{T}^1 .

Extend Σ to a triangulation Σ' on B by coning each simplex of Σ to the point z . Suppose $\{D_0, \dots, D_{n-1}\}$ is the image of an $(n-1)$ -simplex Δ of Σ under the map ι . We now identify the n -simplex of Σ' which is the cone on Δ with the unit cube in \mathbb{R}^n . Label the axes of \mathbb{R}^n with the variables t_0, \dots, t_{n-1} . Place z at the origin, and the vertex v of Δ such that $\iota(v) = D_i$ at the point with $t_i = 1$ and $t_j = 0$ for all $j \neq i$. If p is at the barycenter of a face σ of Δ then place it at the vertex of the cube where the coordinates corresponding to the vertices of σ are 1 and the other coordinates are 0. We now linearly extend over the entire simplex to complete the identification with the cube. Now, if x is in this n -simplex then x has coordinates $(t_0(x), \dots, t_{n-1}(x))$. Let H_x be the surface obtained from H by replacing $H \cap N_i$ with the surface $H_i(t_i(x))$, for each i between 0 and $n-1$. See Figure 13. Repeating this for each n -simplex of Σ' gives us the complete family of surfaces H_x .

We assume H is initially transverse to \mathcal{T}^2 . (That is, small perturbations of H do not change $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$.) For each i , the surface $H_i(t) \subset N_i$ is tangent to \mathcal{T}^2 for finitely many values $\{t_i^j\}$ of t . Hence, for each $x \in B$ the surface H_x is tangent to \mathcal{T}^2 at finitely points, and each such point is in a distinct ball N_i . Note also that if $t_i(x) = t_i(y)$, then H_x and H_y agree inside of N_i . Hence, if H_x is tangent to \mathcal{T}^2 in N_i then the surface H_y will also be tangent to \mathcal{T}^2 , for all y in the plane where $t_i(y) = t_i(x)$. It follows that each n -simplex of Σ' is cubed by the points x where H_x is tangent to \mathcal{T}^2 . See Figure 14. Hence, B is cubed by the n -simplices of Σ' , together with this cubing of each such simplex. We denote this cubing of B as Ω . It follows that if x is in a codimension p cell of Ω then the surface H_x is tangent to \mathcal{T}^2 in at most p points.

We now produce a contradiction by defining a continuous map Ψ from B into $[\mathcal{CE}(H, \mathcal{T}^1)]$ which maps $S = \partial B$ onto $\iota(S)$. The map $\Psi|_S$ will be equal to ι on the barycenters of the $(n-1)$ -cells of Σ , which




FIGURE 13. A simplex Δ of Σ' , and a few of the surfaces H_x for $x \in \Delta$. The union of the faces of the cube that do not meet z is a simplex of T .




FIGURE 14. A simplex Δ of Σ' is cut up by planes into subcubes. Each such plane is determined by the points x in which H_x is tangent to \mathcal{T}^2 in N_i , for some i .

will in turn imply that Ψ maps S onto $\iota(S)$ with the same degree as ι . A contradiction follows as $\iota(S)$ is not homotopic to a point.

For each $x \in B$ let $V_x = [\mathcal{CE}_{\mathcal{T}^2}(H_x, \mathcal{T}^1)]$. If τ is a cell of Ω , then we define V_τ to be the set V_x , for any choice of x in the interior of τ . Note that if x and y are in the interior of the same cell τ of Ω , then the pair (H_x, \mathcal{T}^2) is isotopic to (H_y, \mathcal{T}^2) . Hence, $V_x = V_y$, and thus V_τ is well defined. The map Ψ defined below will take each cell τ of Ω into V_τ . First, we establish a few properties of V_τ .

Claim 6.3. *Suppose σ is a cell of Ω which lies on the boundary of a cell τ . Then $V_\sigma \subset V_\tau$.*

Proof. Pick $x \in \sigma$ and $y \in \tau$. If $D \in V_x$ then D is some kind of compression for H_x that can be isotoped to be disjoint from \mathcal{T}^2 . To show $D \in V_y$ we must show that D is some kind of compression for H_y disjoint from \mathcal{T}^2 . Note that $H_y \cap \mathcal{T}^2$ is obtained from $H_x \cap \mathcal{T}^2$ by resolving some tangency. Hence, any loop or arc of $H_x \setminus \mathcal{T}^2$ is isotopic to a loop or arc of $H_y \setminus \mathcal{T}^2$. It follows that since $D \cap H_x$ was a collection of loops and arcs on H_x disjoint from \mathcal{T}^2 , then $D \cap H_y$ will be a similar such collection. Hence, $D \in [\mathcal{CE}_{\mathcal{T}^2}(H_y, \mathcal{T}^1)] = V_y$. \square

Claim 6.4. *For each cell τ of Ω ,*

$$\pi_i(V_\tau) = 1 \text{ for all } i \leq \dim(\tau) - 1.$$

Proof. Let x be in the interior of a codimension p cell τ of Ω . Then the dimension $\dim(\tau)$ is $n - p$. The surface H_x is tangent to \mathcal{T}^2 in at most p points, and is transverse to H_x elsewhere. Recall $V_x = [\mathcal{CE}_{\mathcal{T}^2}(H_x, \mathcal{T}^1)]$. Thus, if the theorem is false then V_x is non-empty, and $\pi_i(V_x) = 1$ for all $i \leq n - p - 1 = \dim(\tau) - 1$. \square

We now define Ψ on the 0-skeleton of Ω . For each 0-cell $x \in \Omega$, we will choose a point in V_x to be $\Psi(x)$. If x is in the interior of B then $\Psi(x)$ may be chosen to be an arbitrary point of V_x . If x is a point of $S = \partial B$ then x is contained in (perhaps more than one) $(n - 1)$ -simplex Δ_x of Σ . Let Δ'_x denote the face of Δ_x spanned by the vertices v such that $t_i(v) = 1$ if $t_i(x) = 1$, and $t_i(v) = 0$ otherwise. (Note that if x was on the boundary of Δ_x , so that it was also contained in some other $(n - 1)$ -simplex of Σ , then we still end up with the same simplex Δ'_x of Σ .) So, for example, if x is at the barycenter of Δ_x then $\Delta'_x = \Delta_x$. By construction, for each vertex v of Δ'_x the surface H_x has a compression of some kind D disjoint from \mathcal{T}^2 . Hence, for all y near x the disk D is a compression of some kind for H_y that is disjoint from \mathcal{T}^2 . It follows that the entire simplex $\iota(\Delta'_x)$ is contained in V_x , and

thus we may choose the barycenter of $\iota(\Delta'_x)$ to be the image of $\Psi(x)$. In particular, if x is the barycenter of Δ_x then $\Psi(x) = \iota(x)$.

We now proceed to define the rest of the map Ψ by induction. Let τ be a d -dimensional cell of Ω . By induction, assume Ψ has been defined on the $(d-1)$ -skeleton of Ω . In particular, Ψ has been defined on $\partial\tau$. Suppose σ is a face of τ . By Claim 6.3 $V_\sigma \subset V_\tau$. By assumption $\Psi|_\sigma$ is defined and $\Psi(\sigma) \subset V_\sigma$. We conclude $\Psi(\sigma) \subset V_\tau$ for all $\sigma \subset \partial\tau$, and thus

$$(3) \quad \Psi(\partial\tau) \subset V_\tau.$$

Since $d = \dim(\tau)$ it follows from Claim 6.4 that $\pi_{(d-1)}(V_\tau) = 1$. Since $d-1$ is the dimension of $\partial\tau$, we can thus extend Ψ to a map from τ into V_τ .

What remains to be shown is that if τ is in $S = \partial B$ then the extension of Ψ from $\partial\tau$ to τ may be done so that $\Psi(\tau) \subset \iota(S)$. Let Δ_τ be the simplex of Σ whose interior contains τ . We need only show that $\Psi(\partial\tau) \subset V_\tau \cap \iota(\Delta_\tau)$. Since $V_\tau \cap \iota(\Delta_\tau)$ will be a subsimplex of $\iota(\Delta_\tau)$, it follows that Ψ can be extended over τ to this subsimplex.

By Equation 3, $\Psi(\partial\tau) \subset V_\tau$. So all we must do now is to show $\Psi(\partial\tau) \subset \iota(\Delta_\tau)$. Let σ denote a face of τ , and Δ_σ the simplex of Σ whose interior contains σ . Then Δ_σ is contained in Δ_τ . By induction we may assume $\Psi(\sigma) \subset \iota(\Delta_\sigma)$. Putting this together we conclude $\Psi(\sigma) \subset \iota(\Delta_\tau)$ for each $\sigma \subset \partial\tau$, and thus $\Psi(\partial\tau) \subset \iota(\Delta_\tau)$. \square

7. STAGE 5: $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)] \rightarrow [\mathcal{CE}(H, \mathcal{T}^2)]$

Definition 7.1. Let $M^1 = M - N(\mathcal{T}^1)$ and $H^1 = H \cap M^1$.

Lemma 7.2. [cf. [Bac10], Lemma 3.6.] *Suppose $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ has well-defined homotopy index. Let $D \in \mathcal{CE}(H, \mathcal{T}^2) - \mathcal{CE}(H, \mathcal{T}^1)$. Then $\mathcal{CE}_{\mathcal{T}^2}(H/D, \mathcal{T}^1) = \mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)$.*

Proof. Note that $\mathcal{E}(H, \mathcal{T}^2) \subset \mathcal{E}_{\mathcal{T}^2}(H, \mathcal{T}^1) \subset \mathcal{E}(H, \mathcal{T}^1)$. Thus, the disk D must be an element of $\mathcal{C}(H, \mathcal{T}^2)$ that is not in $\mathcal{C}(H, \mathcal{T}^1)$. We conclude ∂D cuts off a subdisk $\mathcal{D} \subset H^1$ that meets \mathcal{T}^2 .

Let $M(H, \mathcal{T}^2)$ and $N(D)$ be as given in Definition 2.6. Then H/D is obtained from H by removing $N(D) \cap H$ from H and replacing it with the frontier D_* of $N(D)$ in $M(H, \mathcal{T}^2)$ (followed by removing any resulting disk or sphere components).

Claim 7.3. $\mathcal{CE}_{\mathcal{T}^2}(H/D, \mathcal{T}^1) \subset \mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)$.

Proof. Suppose $E \in \mathcal{CE}_{\mathcal{T}^2}(H/D, \mathcal{T}^1)$. Then ∂E can be isotoped off of D_* . If E now meets the ball $N(D)$ then it can be further isotoped so that $E \cap N(D)$ is a collection of disks parallel to D . But then each




FIGURE 15. Since D is not a compressing disk for H^1 , any compressing disk E for H^1/D (right figure) is always isotopic to a compressing disk for H^1 (left figure). If E' is a shadow for E as a compressing disk for H^1/D (right figure), then E' is a shadow for E as a compressing disk for H^1 (left figure).

component of $E \cap N(D)$ can be swapped with a disk parallel to \mathcal{D} . The resulting disk has the same boundary as E , but is disjoint from H . By the irreducibility of M^1 this disk must therefore be properly isotopic to E . See Figure 15. We conclude that E was a compressing or edge-compressing disk for H^1 that persisted as a compressing or edge-compressing disk for H^1/D . E is therefore a compressing or edge-compressing disk for H^1 that is disjoint from D .

Now let E' be a shadow for E as a compressing or edge-compressing disk for H^1/D . As $\partial E' = \partial E$, it follows that $\partial E' \cap D_* = \emptyset$. So, if E' meets the ball $N(D)$, then it meets it in disks parallel to D . The disk E' thus meets H in loops/arcs isotopic to $E' \cap H/D$, together with loops/arcs parallel to $D \cap H$. It follows that the interior of E' meets H in inessential loops/arcs, and thus E' is a shadow for E as a compressing or edge-compressing disk for H^1 , i.e. $E \in \mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)$. See Figure 15. \square

Claim 7.4. $\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1) \subset \mathcal{CE}_{\mathcal{T}^2}(H/D, \mathcal{T}^1)$.

Proof. Let E now denote an element of $\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)$. Thus, $\partial E \cap \mathcal{T}^2 = \emptyset$. We assume E has been chosen so that $|E \cap \mathcal{D}|$ is minimal. First we suppose $E \cap \mathcal{D} = \emptyset$. If the interior of E meets D then we may surger it off by a standard innermost disk argument. So in this case we may assume $E \cap D = \emptyset$. Since E is a compressing or edge-compressing disk for H^1 but D is not, it now follows that E is a compressing or edge-compressing disk for H^1/D . Any shadow for E as a compressing or edge-compressing disk for H^1 will be a shadow for E as a compressing or edge-compressing disk for H^1/D , and thus $E \in [\mathcal{CE}_{\mathcal{T}^2}(H/D, \mathcal{T}^1)]$.




FIGURE 16. Using the disk \mathcal{D}' to obtain E_0 and E'_0 from E and E' .

Finally, we consider the case $E \cap \mathcal{D} \neq \emptyset$. Our goal is to isotope E to a compressing or edge-compressing disk $E_0 \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ such that $|E_0 \cap \mathcal{D}| < |E \cap \mathcal{D}|$, contradicting our minimality assumption.

Let γ denote an arc of $\partial E \cap \mathcal{D}$ that is outermost on \mathcal{D} . Then γ cuts a disk \mathcal{D}' off of \mathcal{D} whose interior does not meet E . We can use the disk \mathcal{D}' to guide an isotopy of both E and its shadow E' to a compressing disk E_0 for H and a disk E'_0 with $\partial E_0 = \partial E'_0$. See Figure 16. Note that $|E_0 \cap \mathcal{D}| < |E \cap \mathcal{D}|$. If $\mathcal{D}' \cap \mathcal{T}^2 = \emptyset$, then it follows from the fact that E' was a shadow of E that E'_0 will be a shadow of E_0 . Thus, $E_0 \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ as desired.

If $\mathcal{D}' \cap \mathcal{T}^2 \neq \emptyset$ then the disk E'_0 will not be a shadow for E_0 , since $E'_0 \cap \mathcal{T}^2 \neq \emptyset$. What remains then is to show that nonetheless, E_0 has a shadow.

Let $N(\mathcal{T}^2)$ denote a small product neighborhood of \mathcal{T}^2 . Since $E'_0 \cap \mathcal{T}^2 \neq \emptyset$, it follows that $E'_0 \cap \partial N(\mathcal{T}^2) \neq \emptyset$. Let δ denote a loop of $E'_0 \cap \partial N(\mathcal{T}^2)$ that is outermost on E'_0 . As \mathcal{T}^2 is incompressible, δ bounds a subdisk F_* of $\partial N(\mathcal{T}^2)$. See Figure 17.

Although \mathcal{T}^2 may not be transverse to H , the surface $\partial N(\mathcal{T}^2)$ will be. Thus, the disk F_* meets H in a collection of loops. We claim these loops are inessential on H^1 , and thus F_* can be used to surger E'_0 to a disk which meets \mathcal{T}^2 fewer times. The new disk will meet H more times, but each new intersection introduced will be inessential on H^1 . Thus, by repeating this process we transform E'_0 to a shadow for E_0 , as desired.

To obtain a contradiction, suppose at least one loop of $F_* \cap H$ is essential on H^1 . Let α be a such loop that is innermost on F_* . The loop α bounds a subdisk A' of F_* whose interior may meet H^1 in inessential




FIGURE 17. The curves α , β , and δ , and the disks A' , B , and F_* .

loops. See Figure 17. We claim A' is the shadow of a compressing disk A for H^1 , and thus $A \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$.

Let β denote a loop of $A' \cap H$ that is innermost on A' . As β is inessential on H^1 it bounds a subdisk B of H^1 . See Figure 17. The disk B can be used to surger A' , lowering $|A' \cap H|$. Continuing in this way we arrive at a disk A with the same boundary as A' but whose interior is disjoint from H . As $\partial A' = \partial A$ is essential on H^1 , we conclude A is a compressing disk for H^1 . The disk A' is then a shadow for A , and thus $A \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$.

Finally, suppose X is any other element of $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$. As $\partial A = \partial A' \subset F_* \subset \partial N(\mathcal{T}^2)$ and $\partial X \cap \mathcal{T}^2 = \emptyset$, it follows that $\partial X \cap \partial A = \emptyset$. By a standard innermost disk argument (and the irreducibility of M^1) we may isotope X to remove any intersections of its interior with the interior of A . Thus, we may assume $A \cap X = \emptyset$. The disk X is therefore connected to the disk A by an edge in $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$. As this holds for all disks $X \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$, we conclude $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ is contractible to A . As $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ is not contractible, we have reached a contradiction. \square

The two claims complete the proof of Lemma 7.2. \square

At the end of the previous section we were left with a surface H with $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ having a well-defined homotopy index. In the next theorem we transition to a similar statement about $[\mathcal{CE}(H, \mathcal{T}^2)]$.

Theorem 7.5. *Suppose H is transverse to \mathcal{T}^2 away from a finite number of points of tangency, and $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ has homotopy index n . Then H may be isotoped so that any point of tangency is a saddle contained in \mathcal{T}^2 , and $[\mathcal{CE}(H, \mathcal{T}^2)]$ has homotopy index n .*

Proof. Suppose $D \in \mathcal{CE}(H, \mathcal{T}^2) - \mathcal{CE}(H, \mathcal{T}^1)$. Then ∂D cuts off a subdisk D' of H^1 that meets \mathcal{T}^2 . The surface H/D is obtained by exchanging D' with D . By the incompressibility of ∂M and the irreducibility of M , this surface is isotopic to H in M . As H/D meets \mathcal{T}^2 fewer times, a similar surgery can only be performed a finite number of times. Thus, after a maximal sequence of such surgeries we obtain a surface isotopic to H (which we continue to denote by the same letter), where every compressing or edge-compressing disk for $H - N(\mathcal{T}^2)$ is a compressing or edge-compressing disk for H^1 . It follows that there are no longer any center tangencies in $H \cap \mathcal{T}^2$. Furthermore, by Lemma 7.2 the complex $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$ has remained unchanged. Thus, if this complex had homotopy index n to begin with, then it still has homotopy index n .

Now suppose $D \in [\mathcal{CE}(H, \mathcal{T}^2)]$. Then by the previous paragraph, $D \in [\mathcal{CE}(H, \mathcal{T}^1)]$. Furthermore, as D lies in the complement of \mathcal{T}^2 , $D \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$. We conclude $[\mathcal{CE}(H, \mathcal{T}^2)] \subset [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$. We now claim the opposite inclusion is true as well, and thus $[\mathcal{CE}(H, \mathcal{T}^2)] = [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$.

Suppose $E \in [\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)]$. Let E' be a shadow of E . Let β be a loop of $E' \cap H^1$ that is innermost on E' . Then β bounds subdisks $C \subset E'$ and $C' \subset H^1$. If $C' \cap \mathcal{T}^2 \neq \emptyset$, then C is a compressing disk for $H - N(\mathcal{T}^2)$ that is not a compressing disk for H^1 , a contradiction. We conclude $C' \cap \mathcal{T}^2 = \emptyset$. Since $E' \cap \mathcal{T}^2 = \emptyset$ and $C \subset E'$, we conclude $C \cap \mathcal{T}^2 = \emptyset$. The sphere $C \cup C'$ thus bounds a ball in the complement of \mathcal{T}^2 that we can use to guide an isotopy of C to C' . (This may remove other components of $E' \cap C$ as well.) We thus transform the disk E' to a disk E'' such that $\partial E'' = \partial E$, $E'' \cap \mathcal{T}^2 = \emptyset$, and $|E'' \cap H^1| < |E' \cap H^1|$. Continuing in this way we arrive at a compressing disk for H^1 with the same boundary as E , which is disjoint from \mathcal{T}^2 . Thus $E \in [\mathcal{CE}(H, \mathcal{T}^2)]$, completing the proof that $[\mathcal{CE}_{\mathcal{T}^2}(H, \mathcal{T}^1)] \subset [\mathcal{CE}(H, \mathcal{T}^2)]$ \square

Lemma 7.6. *Suppose H is transverse to \mathcal{T}^2 away from a finite number of saddle tangencies, and $[\mathcal{CE}(H, \mathcal{T}^2)]$ has well-defined homotopy index. Then H is transverse to \mathcal{T}^2 , and meets each 2-simplex of \mathcal{T}^2 in a collection of normal arcs.*

Proof. Suppose Δ is a 2-simplex of \mathcal{T}^2 . If $H \cap \Delta$ contains a loop, a non-normal arc, or a saddle tangency, then some subdisk D of Δ that can be pushed off \mathcal{T}^2 to form a compressing or edge-compressing

disk for $H - N(\mathcal{T}^2)$. We conclude $[\mathcal{CE}(H, \mathcal{T}^2)]$ could not have been empty. As D is isotopic into $\partial(M - N(\mathcal{T}^2))$, it can be made disjoint from any other element of $[\mathcal{CE}(H, \mathcal{T}^2)]$. This implies every vertex of $[\mathcal{CE}(H, \mathcal{T}^2)]$ is connected to D by an edge, contradicting the assumption that $[\mathcal{CE}(H, \mathcal{T}^2)]$ is not contractible. \square

8. FROM H TO THE COMPONENTS OF $H - N(\mathcal{T}^2)$.

Theorem 8.1. [cf. [Bac10], Theorem 4.7.] *Suppose $\text{ind}[\mathcal{CE}(H, \mathcal{T}^2)] = n$. Then*

$$\sum_{\Delta \in \mathcal{T}^3} \text{ind}[\mathcal{CE}(H \cap \Delta, \mathcal{T}^2)] = n.$$

Proof. Let $\{\Delta_i\}$ denote the 3-simplices in \mathcal{T}^3 . Let $H_i = H \cap \Delta_i^3$. Notice that elements of $\mathcal{CE}(H_i, \mathcal{T}^2)$ and $\mathcal{CE}(H_j, \mathcal{T}^2)$ are disjoint when $i \neq j$. Hence, $[\mathcal{CE}(H, \mathcal{T}^2)]$ is the join of all of the complexes $[\mathcal{CE}(H_i, \mathcal{T}^2)]$.

The proof is by induction on m , the number of tetrahedra in \mathcal{T} . The base case $m = 1$ is trivial. Let $G = \bigcup_{i=2}^m H_i$, so that $[\mathcal{CE}(H, \mathcal{T}^2)]$ is the join of $[\mathcal{CE}(H_1, \mathcal{T}^2)]$ and $[\mathcal{CE}(G, \mathcal{T}^2)]$.

We first observe that $[\mathcal{CE}(H_1, \mathcal{T}^2)]$ and $[\mathcal{CE}(G, \mathcal{T}^2)]$ are not contractible. This is because the join of a contractible space with any other space is also contractible. We conclude that if either complex is contractible then the complex $[\mathcal{CE}(H, \mathcal{T}^2)]$ would have been contractible, a contradiction. Since neither is contractible, each has a well-defined homotopy index.

If either H_1 or G has local index 0 then the result is immediate, as the join of a complex with the empty set is the same complex. We assume, then, that the local index of H_1 is $n \geq 1$ and the local index of G is $m \geq 1$.

By definition, $(n-1)$ is the smallest i such that $\pi_i([\mathcal{CE}(H_1, \mathcal{T}^2)]) \neq 1$, and $(m-1)$ is the smallest j such that $\pi_j([\mathcal{CE}(G, \mathcal{T}^2)]) \neq 1$. Our goal is to show that $(n+m-1)$ is the smallest k such that

$$\pi_k([\mathcal{CE}(H_1 \cup G, \mathcal{T}^2)]) = \pi_k([\mathcal{CE}(H_1, \mathcal{T}^2)] * [\mathcal{CE}(G, \mathcal{T}^2)]) \neq 1.$$

When $n = 2$ then $\pi_1([\mathcal{CE}(H_1, \mathcal{T}^2)]) \neq 1$. Suppose H_1 separates Δ_1^3 into \mathcal{V} and \mathcal{W} . Let $\mathcal{V}(H_1)$ and $\mathcal{W}(H_1)$ denote the subsets of $\mathcal{CE}(H_1, \mathcal{T}^2)$ spanned by the compressing and edge-compressing disks that lie in \mathcal{V} and \mathcal{W} , respectively. By an argument identical to the one given by McCullough in [McC91], $[\mathcal{V}(H_1)]$ and $[\mathcal{W}(H_1)]$ are contractible. If we contract these to points $p_{\mathcal{V}}$ and $p_{\mathcal{W}}$, then the remaining 1-simplices of $[\mathcal{CE}(H_1, \mathcal{T}^2)]$ join these two points. The fundamental group $\pi_1([\mathcal{CE}(H_1, \mathcal{T}^2)])$ is generated by these 1-simplices.

The remaining 2-simplices have become bigons that run once over each of two 1-simplices. Hence, each such 2-simplex gives rise to a relation in $\pi_1([\mathcal{CE}(H_1, \mathcal{T}^2)])$ that kills one generator. It follows that $\pi_1([\mathcal{CE}(H_1, \mathcal{T}^2)])$ is free, and hence the non-triviality of $\pi_1([\mathcal{CE}(H_1, \mathcal{T}^2)])$ implies that the homology group $H_1([\mathcal{CE}(H_1, \mathcal{T}^2)])$ is also non-trivial. Similarly, if $m = 2$ we conclude $H_1([\mathcal{CE}(G, \mathcal{T}^2)])$ is non-trivial. For $n \geq 3$ the non-triviality of $H_{n-1}([\mathcal{CE}(H_1, \mathcal{T}^2)])$ follows from the Hurewicz Theorem.

By Lemma 2.1 from [Mil68]:

$$\begin{aligned} \tilde{H}_{n+m-1}([\mathcal{CE}(H_1, \mathcal{T}^2)] * [\mathcal{CE}(G, \mathcal{T}^2)]) \\ \cong \sum_{i+j=n+m-2} \tilde{H}_i([\mathcal{CE}(H_1, \mathcal{T}^2)]) \otimes \tilde{H}_j([\mathcal{CE}(G, \mathcal{T}^2)]) \\ + \sum_{i+j=n+m-3} \text{Tor}(\tilde{H}_i([\mathcal{CE}(H_1, \mathcal{T}^2)]), \tilde{H}_j([\mathcal{CE}(G, \mathcal{T}^2)])). \end{aligned}$$

In particular, it follows from the fact that $(n-1)$ is the smallest i such that $H_i([\mathcal{CE}(H_1, \mathcal{T}^2)])$ is non-trivial, and $(m-1)$ is the smallest j such that $H_j([\mathcal{CE}(G, \mathcal{T}^2)])$ is non-trivial, that $(n+m-1)$ is the smallest k such that $H_k([\mathcal{CE}(H_1, \mathcal{T}^2)] * [\mathcal{CE}(G, \mathcal{T}^2)])$ is non-trivial. □

REFERENCES

- [Baca] D. Bachman. 2-normal surfaces. Preprint. Available at <http://arxiv.org/abs/math/0309437>.
- [Bacb] D. Bachman. Normalizing Topologically Minimal Surfaces II: Disks. Preprint.
- [Bacc] D. Bachman. Normalizing Topologically Minimal Surfaces III: Index 2 Surfaces. Preprint.
- [Bacd] David Bachman. Stabilizing and destabilizing Heegaard splittings of sufficiently complicated 3-manifolds. *Mathematische Annalen*, pages 1–32. 10.1007/s00208-012-0802-4.
- [Bac02] D. Bachman. Critical Heegaard surfaces. *Trans. Amer. Math. Soc.*, 354(10):4015–4042 (electronic), 2002.
- [Bac10] D. Bachman. Topological Index Theory for surfaces in 3-manifolds. *Geometry & Topology*, 14(1):585–609, 2010.
- [BDTSa] D. Bachman, R. Derby-Talbot, and E. Sedgwick. Almost Normal Surfaces with boundary. Preprint. Available at <http://arxiv.org/abs/1203.4632>.
- [BDTSb] D. Bachman, R. Derby-Talbot, and E. Sedgwick. Most Dehn fillings preserve the set of Heegaard splittings. In preparation.
- [CG87] A. J. Casson and C. McA. Gordon. Reducing Heegaard splittings. *Topology and its Applications*, 27:275–283, 1987.

- [CM04a] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. I. Estimates off the axis for disks. *Ann. of Math. (2)*, 160(1):27–68, 2004.
- [CM04b] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. II. Multi-valued graphs in disks. *Ann. of Math. (2)*, 160(1):69–92, 2004.
- [CM04c] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. III. Planar domains. *Ann. of Math. (2)*, 160(2):523–572, 2004.
- [CM04d] Tobias H. Colding and William P. Minicozzi, II. The space of embedded minimal surfaces of fixed genus in a 3-manifold. IV. Locally simply connected. *Ann. of Math. (2)*, 160(2):573–615, 2004.
- [Hak61] W. Haken. Theorie der Normalfächen. *Acta Math.*, 105:245–375, 1961.
- [Hak68] W. Haken. *Some Results on Surfaces in 3-Manifolds*. M.A.A., Prentice Hall, 1968.
- [Hem01] J. Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40:631–657, 2001.
- [Joh] J. Johnson. Calculating Isotopy Classes of Heegaard Splittings. Preprint. Available at <http://arxiv.org/abs/1004.4669>.
- [Kne29] H. Kneser. Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten. *Jahresbericht der Dent. Math. Verein*, 28:248–260, 1929.
- [McC91] Darryl McCullough. Virtually geometrically finite mapping class groups of 3-manifolds. *J. Differential Geom.*, 33(1):1–65, 1991.
- [Mil68] J. Milnor. *Morse Theory*. PUP, Princeton, New Jersey, 1968.
- [Rub95] J. H. Rubinstein. Polyhedral minimal surfaces, Heegaard splittings and decision problems for 3-dimensional manifolds. In *Proceedings of the Georgia Topology Conference*, pages 1–20, 1995.
- [Sto00] M. Stocking. Almost normal surfaces in 3-manifolds. *Trans. Amer. Math. Soc.*, 352:171–207, 2000.